Fermilab

Title: Experience with DØ Data Handling in Production

We report on the production experience of the DØ experiment at the Fermilab Tevatron, using the SAM data handling system with a variety of computing hardware configurations, batch systems, and mass storage strategies. We have stored metadata representing more than 250 TB of data, of which 220 TB is physically stored in the Fermilab ENSTORE mass storage system and 30 TB is virtual data. [These numbers are captured on December 1, 2002, and continue to grow at a rate of more than 1 TB per day.] We deliver data through this system at an average rate of more than 2 TB/day to analysis programs on the largest SAM station, with a substantial multiplication factor in the consumed data due to user interest in the same samples. We handle more than one million files in this system. We have provided data delivery to user jobs at Fermilab on three different types of systems: a large SMP system, a Linux batch cluster, and a Linux desktop cluster. In addition, we import simulation data generated at 6 sites worldwide, and deliver data to jobs at many more sites. We describe the scope of the data handling deployment worldwide, the operational experience with this system and the feedback of that experience into design changes, monitoring requirements, and desirable optimizations. We also describe the analysis strategies supported by this system.

-Baranovski, A.; Bonham, D.; Brock, R.; Carpenter, L.; Lueking, L., Merritt, W., Terekhov, I.; Trumbo, J.; Veseli, S.; Weigand, J.; White, S.; Yip, K.